General-purpose 128 × 128 SIMD processor array with integrated image sensor

P. Dudek and S.J. Carey

A CMOS image sensor/processor chip fabricated in a 0.35 μm CMOS technology is presented. The chip contains a general purpose software-programmable SIMD array of 128 × 128 processing elements. It executes over 20 GOPS while dissipating 240 mW of power and achieves pixel-processor density of 410 cells/mm^2. Performance and accuracy measurement results are given.

Introduction: In some computer vision applications, especially where high computational performance is required together with low power consumption, it is beneficial to use 'vision chips'—devices that combine image sensing and processing on a single silicon die. Application-specific vision chips can efficiently perform some uncomplicated low-level image processing tasks, such as filtering, edge detection, or motion detection, often using analogue circuitry operating in pixel-parallel fashion. But to combine different operations to form a more complex algorithm, and to provide flexibility to use a chip for various algorithms, a software-programmable solution must be employed. The required flexibility can be provided by general-purpose SIMD processor arrays. We have been working towards combining the two paradigms: a general-purpose SIMD array and analogue processor-per-pixel vision systems [1–3]. Our first proof-of-concept 'analogue SIMD vision chip' [1] (SCAMP—switched current analogue matrix processor), was a 21 × 21 processor array. An improved, 39 × 48 chip version (SCAMP-2) was subsequently developed [2] followed recently by a 128 × 128 (SCAMP-3) device [3]. In this Letter we overview the SCAMP-3 chip design and present measured performance/accuracy data and experimental image processing results.

Architecture: The architecture of the chip is shown in Fig. 1. The chip contains a 128 × 128 array of analogue processing elements (APEs). While these processors are implemented using analogue circuitry, their architecture and functionality are similar to simple digital PEs—they execute instruction-level programs, performing arithmetic and logic operations on data stored in the local memory. Each APE includes nine 'registers' (one of them doubling as a neighbour communication port), which can store analogue samples of data. APEs can communicate and exchange data with four nearest neighbours. All APEs execute identical instructions, which are broadcast by a single controller. Local autonomy is provided via activity-flag registers, which can disable selected APEs conditionally, depending on local data. Each APE also contains a photodetector circuit.

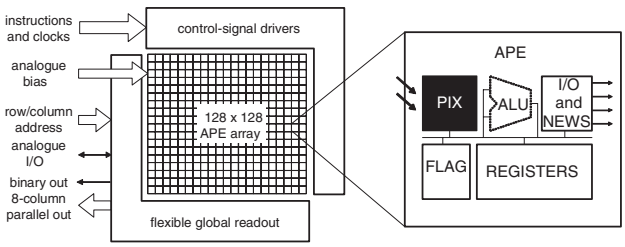

Fig. 1 *Architecture of SCAMP-3 vision chip*

The APE circuit is implemented using switched-current memory cells. Each APE is capable of register-transfer and arithmetic operations, which include inversion, addition and division by two, with the important design feature that no actual hardware circuitry (in addition to registers and bus) is required to implement the arithmetic operations. The operation of the processor follows the principles described in [1]; however, the multiplier has been replaced by current division, which is performed by reading the current from one register and storing it to two registers at the same time.

Implementation: The chip has been implemented in a 0.35 μm 1-poly 3-metal CMOS technology (Austria Micro Systems). The chip operates with 3.3 V (digital) and 2.5 V (analogue) supply voltages. An APE (including photodetector) occupies a silicon area of 49.35 × 49.35 μm. Control signals, analogue bias voltages and power supply lines have been routed over the APE area. The pixel fill-factor is approximately 5.6%.

To operate the chip requires a number of analogue bias voltages, provided and adjusted externally, plus a sequence of digital instruction codewords, which are provided from an external sequencer/controller (currently implemented on FPGA). Control signals are generated on chip, combining instruction-codewords with clock phases, and routed to each APE using horizontal and vertical wires, with one buffer/driver for each row and column of the array.

The processing results are read-out from the array via flexible global readout circuitry, which enables analogue, binary and column parallel (8-bit) readout. The readout architecture also implements calculation of global sum and logic OR operations, on the entire array or parts of the array, which is useful in many algorithms (e.g. pixel counts, loop iteration control, extraction of object co-ordinates, histograms etc.). The output interface allows analogue readout at a maximum speed of 3 Mpixel/s; faster readout speeds are available for binary images.

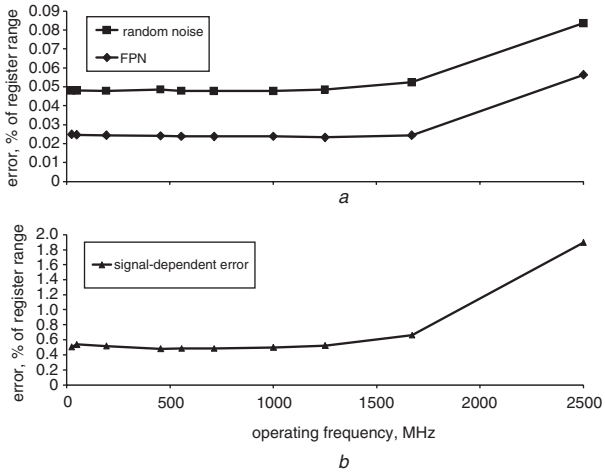

Fig. 2 *Measured error of register transfer operation against operating frequency*

a Random noise and fixed pattern noise
b Signal-dependent error

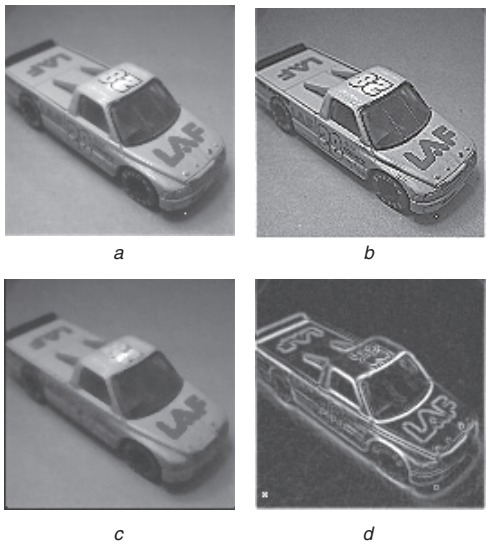

Fig. 3 *Examples of image processing executed on chip*

a Acquired image
b Sharpening filter
c 3 × 3 median filter
d Sobel edge detection

Results: An important aspect of the analogue processor design is the accuracy of storage and arithmetic operations. The measured signal-dependent error of the register-transfer was 0.52% of full scale range. The noise associated with the transfer operation was 0.048% RMS of

the register full scale range. The fixed pattern error (calculated over the entire array) imposed upon register transfer operations was 0.024% RMS. It should be noted that in practice it is the fixed-pattern-error which is the most important source of distortion in image processing using analogue cellular arrays [4]. The above measurements were performed with a 1 MHz clock. The chip operates consistently with an instruction clock up to 1.25 MHz; higher frequency of operation is possible but with degraded accuracy. The increase of the errors with frequency is illustrated in Fig. 2. The measured accuracy of the 'division-by-two' operation (including algorithmic error compensation) was 0.12% RMS of register full scale range (average fixed pattern error, measured with a 1 MHz clock).

Fig. 3 shows results of performing some typical low-level image processing on the chip. Power consumption depends on the frame rate and the program being executed. For simple operations at typical frame rates, power can be as low as a few mW. For example, when performing Sobel edge detection at 30 frames per second, with binary readout (thresholded result of the grey-scale edge map), the chip dissipates below 2 mW. Operating at the maximum computational power (20 GOPS with a 1.25 MHz clock) the chip dissipates 240 mW.

Acknowledgment: This work has been supported by grants from the Engineering and Physical Sciences Research Council (EPSRC).

3 February 2006
Electronics Letters online no: 20060352
doi: 10.1049/el:20060352

P. Dudek and S.J. Carey (*School of Electrical and Electronic Engineering, The University of Manchester, PO Box 88, Manchester M60 1QD, United Kingdom*)

E-mail: pdudek@umist.ac.uk

References

1 Dudek, P., and Hicks, P.J.: 'A general-purpose processor-per-pixel analog SIMD vision chip', *IEEE Trans. Circuits and Systems—I: Regular papers*, 2005, **52**, (1), pp. 13–20
2 Dudek, P.: 'A 39×48 general-purpose focal-plane processor array integrated circuit'. Proc. ISCAS'04, May 2004, Vol. V, pp. 449–452
3 Dudek, P.: 'Implementation of SIMD vision chip with 128×128 array of analogue processing elements'. Proc. ISCAS'05, May 2005, pp. 5806–5809
4 Dudek, P.: 'Accuracy and efficiency of gray-level image filtering on VLSI cellular processor arrays'. Proc. CNNA'04, June 2004, pp. 123–128